

Multi-level methods and approximating distribution functions

D. Wilson^{1, a)} and R.E. Baker¹

¹⁾*Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter,
Woodstock Road, Oxford, OX2 6GG*

Biochemical reaction networks are often modelled using discrete-state, continuous-time Markov chains. System statistics of these Markov chains usually cannot be calculated analytically and therefore estimates must be generated via simulation techniques. There is a well documented class of simulation techniques known as exact stochastic simulation algorithms, an example of which is Gillespie’s direct method. These algorithms often come with high computational costs, therefore approximate stochastic simulation algorithms such as the tau-leap method are used. However, in order to minimise the bias in the estimates generated using them, a relatively small value of tau is needed, rendering the computational costs comparable to Gillespie’s direct method.

The multi-level Monte Carlo method (Anderson and Higham, Multiscale Model. Simul. 10:146–179, 2012) provides a reduction in computational costs whilst minimising or even eliminating the bias in the estimates of system statistics. This is achieved by first crudely approximating required statistics with many sample paths of low accuracy. Then correction terms are added until a required level of accuracy is reached. Recent literature has primarily focussed on implementing the multi-level method efficiently to estimate a single system statistic. However, it is clearly also of interest to be able to approximate entire probability distributions of species counts. We present two novel methods that combine known techniques for distribution reconstruction with the multi-level method. We demonstrate the potential of our methods using a number of examples.

^{a)}Electronic mail: daniel.wilson@dtc.ox.ac.uk

I. INTRODUCTION

Biochemical reaction networks describe interactions between the molecular populations of biological and physiological systems. A classical deterministic approach to modelling these networks is to use systems of ordinary or partial differential equations to describe the evolution of the concentrations of each species. However, experimental researchers, such as Elowitz et al.², have shown that stochasticity can be observed in a variety of biological phenomena including gene expression within the cell. Thus, deterministic modelling can neglect important characteristics of a system. For example, in systems with low molecular populations using a deterministic model to describe the molecular concentrations can fail to predict a phenomenon known as stochastic focussing in signalling networks²⁰. Deterministic approaches can also fail to accurately replicate the dynamics of systems with large molecular numbers, such as systems operating near bifurcation points⁴.

To include stochasticity into models of biochemical reaction networks one can consider the framework of the chemical master equation¹⁰. This assumes spatial homogeneity in the system, which is the case we shall consider in this paper. The chemical master equation describes the temporal evolution of molecular species in terms of the number of molecules present rather than their concentrations. This description provides a system of ordinary differential equations (ODEs, one for each possible state of the system) that, for simple cases involving only zeroth and first order reactions, can be solved analytically¹⁰. For more complicated systems that include higher order reactions analytical solutions are not generally feasible so one may try and use a numerical scheme to approximate the solution to the chemical master equation^{3,9,10}. However, this is also often infeasible due to the high (possibly infinite) number of distinct states in the system, so the only option to explore system behaviour is the use of simulation techniques.

There is a class of simulation techniques known as exact stochastic simulation algorithms (eSSAs) which model systems precisely, in the sense that the sample paths generated using them are mathematically equivalent to the chemical master equation. However, the more complex the system, the longer these simulations take to run, and for practical purposes approximate stochastic simulation algorithms (aSSAs) that favour speed over accuracy are often used. Unfortunately, obtaining accurate summary statistics for a system using aSSAs

often renders the computational costs of those algorithms on a par with those of eSSAs.

Recently proposed by Anderson and Higham¹ is the discrete-state multi-level method. This method vastly reduces the computational cost associated with generating summary statistics for a given system compared with traditional eSSAs and aSSAs. It does so by generating many sample paths of poor accuracy with very little computational effort and then reduces the error of summary statistics generated using these sample paths by adding correction terms from a smaller number of sample paths of higher accuracy. The paper by Lester et al.¹³ provides a clear introduction to the multi-level method with the focus being on implementation. In this work, we show how to exploit the multi-level method to accurately and efficiently estimate cumulative distribution functions (CDFs) of molecular species at a terminal time, T .

In Section II we briefly describe the basic simulation techniques that will be used throughout the paper, namely Gillespie’s direct method⁷ (Gillespie’s DM) and the tau-leap method⁸. We also introduce the multi-level method, stating how to optimally choose the number of sample paths on each level and simulate these paths with low sample variance. Section III focusses on two methods of approximating the CDF of a chosen molecular species at time T , where we use the Schlögl system¹⁶ and a dimerisation model¹ as example cases. Then we combine these two approximation methods with the multi-level method. We investigate how to optimise implementation of our proposed methods, resulting in two adaptive algorithms designed to use the multi-level method efficiently without any prior knowledge of the system. We conclude, in Section IV, with a short discussion.

II. MODELLING AND SIMULATING BIOCHEMICAL REACTION NETWORKS

Let S_1, \dots, S_N represent N distinct populations or species that may be involved in M reaction channels, R_1, \dots, R_M , in a system with volume ν . We assume that the system is well stirred in order that we may ignore any spatial effects in our model. Denote the size of population S_i at time t as $X_i(t)$. We define the state space vector at time t as

$$\mathbf{X}(t) = (X_1(t), \dots, X_N(t))^T. \quad (1)$$

We also introduce the stoichiometric or state-change matrix, $\boldsymbol{\nu} = \{\nu_{ij}\}$, where ν_{ij} is the

change in the copy number of S_i after the reaction channel R_j fires. Thus $\boldsymbol{\nu}$ is an $N \times M$ matrix. If reaction R_j fires at time $t + s$ given no other reactions in $[t, t + s)$, we update the state space vector as follows

$$\mathbf{X}(t + s) = \mathbf{X}(t) + \boldsymbol{\nu}_j, \quad (2)$$

where $\boldsymbol{\nu}_j$ is the stoichiometric vector, defined as the j^{th} column of the stoichiometric matrix,

$$\boldsymbol{\nu}_j = (\nu_{1j}, \dots, \nu_{Nj})^T. \quad (3)$$

Each reaction R_j has a propensity function $a_j(\mathbf{X}(t))$ at time t , this quantity determines the rate at which the j^{th} reaction occurs. The propensity functions, $a_j(\mathbf{X}(t))$, can be calculated by considering the number of ways in which a reaction can occur given the copy numbers at time t . A detailed description can be found in the paper by Erban et al.⁵.

The Kurtz representation¹² allows us to describe evolution of the state space vector as follows:

$$\mathbf{X}(t) = \mathbf{X}(0) + \sum_{j=1}^M Y_1^j \left(\int_0^t a_j(\mathbf{X}(s)) \, ds \right) \cdot \boldsymbol{\nu}_j, \quad (4)$$

where the Y_1^j , $j = 1, \dots, M$, are Poisson processes of unit rate.

A. Gillespie's direct method

Gillespie's DM⁷ is a standard example of an eSSA for generating sample paths from (4). We include a step by step algorithm below:

1. Initialise the copy numbers, $\mathbf{X}(0)$, and the stoichiometric matrix, $\boldsymbol{\nu}$. Choose a terminal time, T , and set $t = 0$.
2. Calculate the propensity functions, $a_j = a_j(\mathbf{X}(t))$, for each reaction channel, j , and $a_0 = \sum_{j=1}^M a_j$.
3. Generate $r_1 \sim \mathbf{U}(0, 1)$ and set $\Delta = 1/a_0 \log(1/r_1)$.
4. If $t + \Delta > T$ then terminate the algorithm, otherwise continue to step 5.
5. Choose the next reaction to occur by finding the minimal k such that $\sum_{j=1}^k a_j > a_0 \times r_2$, where $r_2 \sim \mathbf{U}(0, 1)$.

6. Update the copy numbers, $\mathbf{X}(t + \Delta) = \mathbf{X}(t) + \boldsymbol{\nu}_k$.
7. Let $t = t + \Delta$ and return to step 2.

The main advantage of this algorithm is that it can be used to produce unbiased estimates of system statistics at a terminal time, T . As an example, for the mean copy number of S_i we generate n sample paths to calculate

$$\mathbb{E}(X_i(T)) = \frac{1}{n} \sum_{r=1}^n X_i^{(r)}(T), \quad (5)$$

where $X_i^{(r)}$ represents the copy number of S_i from the r^{th} sample path. The disadvantage of this eSSA is that it simulates every reaction individually, which can render it very computationally expensive. In order to reduce computation time, we now introduce a more efficient algorithm, known as the tau-leap method, which provides sample paths that approximately follow (4).

B. The tau-leap method

The tau-leap method is an aSSA, first proposed by Gillespie⁸. It is designed to reduce computation time by temporarily fixing the propensity functions and firing several reactions during each time interval of length τ . The method is equivalent to taking a Forward Euler approximation¹⁹ to equation (4),

$$\mathbf{X}(t + \tau) = \mathbf{X}(t) + \sum_{j=1}^M Y_1^j(a_j(\mathbf{X}(t)) \cdot \tau) \cdot \boldsymbol{\nu}_j. \quad (6)$$

A step by step algorithm is again presented below:

1. Initialise the copy numbers, $\mathbf{X}(0)$, and the stoichiometric matrix, $\boldsymbol{\nu}$. Choose a terminal time, T , and a time step, τ , such that T/τ is an integer. Set $t = 0$.
2. If $t + \tau > T$ then terminate the algorithm, otherwise continue to step 3.
3. Calculate the propensity functions, $a_j = a_j(\mathbf{X}(t))$, for each reaction channel, j .
4. Generate Poisson random variates, p_j , from $Y_1^j(a_j \cdot \tau)$ for $j = 1, \dots, M$.

5. Update the copy numbers $\mathbf{X}(t + \tau) = \mathbf{X}(t) + \sum_{j=1}^M p_j \cdot \boldsymbol{\nu}_j$.
6. Let $t = t + \tau$ and return to step 2.

An attractive feature of the algorithm is that the accuracy of the approximation is determined by the choice of τ . The smaller the time step, τ , the greater the accuracy of the method, and indeed as $\tau \downarrow 0$ the method is equivalent to Gillespie's DM⁸. However, often the value of τ needed to gain the required level of accuracy is small enough that the CPU time becomes comparable to that of Gillespie's DM, and so for further computational savings we utilise the multi-level Monte Carlo method^{1,6}.

C. The multi-level Monte Carlo method

In this section we provide a very brief introduction to the multi-level Monte Carlo method^{1,6}. Heuristically, the method combines the best attributes of Gillespie's DM and the tau-leap algorithm, to produce fast and unbiased estimates of system statistics. Further information can be found in the work by Anderson and Higham¹ and Giles⁶, as well as Lester et al.¹³.

Suppose we are interested in calculating an estimate of $X_i(T)$. We start at the base level, otherwise known as level 0. The estimate at this level is calculated using the tau-leap algorithm, as discussed in the previous section, with $\tau = \tau_0$. If n_0 samples are generated, then we have the biased estimate,

$$Q_0 = \mathbb{E}(Z_{\tau_0}) = \frac{1}{n_0} \sum_{r=1}^{n_0} Z_{\tau_0}^{(r)}, \quad (7)$$

where $Z_{\tau_0}^{(r)}$ is the copy number of population i at time T from the r^{th} sample using the tau-leap algorithm with $\tau = \tau_0$. Typically, a large τ_0 is chosen so that Q_0 can be calculated quickly. The purpose of the other levels is to reduce the bias of this crude approximation, Q_0 , by adding correction terms.

The first level correction is calculated by generating two sets of n_1 sample paths using the tau-leap aSSA. One set has $\tau = \tau_0$ and the other $\tau = \tau_1 < \tau_0$. From the two sets of paths, we calculate the correction term as follows:

$$Q_1 = \mathbb{E}(Z_{\tau_1} - Z_{\tau_0}) \approx \frac{1}{n_1} \sum_{r=1}^{n_1} (Z_{\tau_1}^{(r)} - Z_{\tau_0}^{(r)}). \quad (8)$$

This correction is then added to the initial estimate from the base level:

$$Q_0 + Q_1 = \mathbb{E}(Z_{\tau_0}) + \mathbb{E}(Z_{\tau_1} - Z_{\tau_0}) = \mathbb{E}(Z_{\tau_1}). \quad (9)$$

This estimator has bias equivalent to just using the tau-leap aSSA with $\tau = \tau_1$. To further decrease the bias of the estimate we introduce time steps $\tau_\ell < \tau_{\ell-1} < \dots < \tau_0$, and then add additional correction levels,

$$Q_\ell = \mathbb{E}(Z_{\tau_\ell} - Z_{\tau_{\ell-1}}) = \frac{1}{n_\ell} \sum_{r=1}^{n_\ell} \left(Z_{\tau_\ell}^{(r)} - Z_{\tau_{\ell-1}}^{(r)} \right), \quad (10)$$

to give

$$\mathbb{Q}_b := \sum_{\ell=0}^L Q_\ell = \mathbb{E}(Z_{\tau_0}) + \sum_{\ell=1}^L \mathbb{E}(Z_{\tau_\ell} - Z_{\tau_{\ell-1}}) = \mathbb{E}(Z_{\tau_L}). \quad (11)$$

This means that the estimator \mathbb{Q}_b has bias equivalent to that generated using the tau-leap aSSA with $\tau = \tau_L$. The number of levels, L , can be chosen such that a desired level of accuracy is reached. However, a final correction level may be used in order to gain an unbiased estimate, if required. Again we generate two sets of paths, one set from a tau-leap algorithm where $\tau = \tau_L$ and one set from an eSSA such as Gillespie's DM, using n_{L+1} samples. We can then calculate

$$Q_{L+1}^\dagger = \mathbb{E}(X'_i - Z_{\tau_L}) = \frac{1}{n_{L+1}} \sum_{r=1}^{n_{L+1}} \left(X_i'^{(r)}(T) - Z_{\tau_L}^{(r)} \right), \quad (12)$$

where X'_i is an identical and independently distributed (i.i.d.) copy of X_i . This last correction can then be added to the previous estimate and, using the linearity of expectation, we have

$$\mathbb{E}(X_i(T)) = \mathbb{E}(Z_{\tau_0}) + \sum_{\ell=1}^L \mathbb{E}(Z_{\tau_\ell} - Z_{\tau_{\ell-1}}) + \mathbb{E}(X'_i - Z_{\tau_L}) = \sum_{\ell=0}^L Q_\ell + Q_{L+1}^\dagger. \quad (13)$$

For the rest of this paper we shall choose $\tau_\ell = \tau_0/K^\ell$ where $K \in \mathbb{N}^+$. This means that the time steps used on consecutive levels are nested and renders the algorithm simple to implement. However, we note that it is not always the most appropriate choice: when dealing with systems that exhibit behaviours on a range of timescales a more sophisticated choice is necessary¹⁴.

The reduction in computational cost for the multi-level method is achieved by introducing two coupling techniques. A tau-leap coupling technique for the sets of sample paths on

the first L levels, and an exact coupling technique for the $(L + 1)^{st}$ level. The coupling is designed to create a positive correlation between coarse, $Z_{\tau_{\ell-1}}^{(r)}$, and fine, $Z_{\tau_{\ell}}^{(r)}$, sample paths so that $\text{Var}(Z_{\tau_{\ell}} - Z_{\tau_{\ell-1}})$ is reduced and fewer sample paths need to be generated to obtain an estimator with a desired accuracy. The following algorithm is the tau-leap coupling technique:

1. Initialise the copy numbers for the coarse and fine paths, $\mathbf{Z}_c(0)$ and $\mathbf{Z}_f(0)$, and the stoichiometric matrix, $\boldsymbol{\nu}$. Choose a terminal time, T , base level leap size, τ_0 , and scaling factor, K . Set $t = 0$.
2. For $\alpha = 1$ to $T/\tau_{\ell-1}$:
 - (a) Calculate the propensity functions at the current time t for the coarse system, $a_j^c = a_j(\mathbf{Z}_c(t))$, for each reaction channel, j .
 - (b) For $\beta = 1$ to K :
 - i. Calculate the propensity functions at the current time t for the fine system, $a_j^f = a_j(\mathbf{Z}_f(t))$, for each reaction channel, j .
 - ii. Define the following three propensity functions:

$$b_j^1 = \min(a_j^c, a_j^f); \quad (14)$$

$$b_j^2 = a_j^c - b_j^1; \quad (15)$$

$$b_j^3 = a_j^f - b_j^1; \quad (16)$$

for each reaction channel, j .

- iii. For $j = 1, \dots, M$ generate Poisson random variates, Y_j^r , each with rate $b_j^r \cdot \tau_{\ell}$ for $r \in \{1, 2, 3\}$.
- iv. Update the copy numbers of the two systems as follows:

$$\mathbf{Z}_c(t + \tau_{\ell}) = \mathbf{Z}_c(t) + \sum_{j=1}^M (Y_j^1 + Y_j^2) \cdot \boldsymbol{\nu}_j; \quad (17)$$

$$\mathbf{Z}_f(t + \tau_{\ell}) = \mathbf{Z}_f(t) + \sum_{j=1}^M (Y_j^1 + Y_j^3) \cdot \boldsymbol{\nu}_j. \quad (18)$$

Let $t = t + \tau_{\ell}$.

Further details, including those detailing the exact coupling technique, can be found in the paper by Lester et al.¹³.

Finally, we need an algorithm to evaluate how many paths, n_ℓ , to generate on each level, ℓ , such that we minimise the expected CPU time of the method whilst controlling the variance, $\mathbb{V} = \sum_{\ell=0}^L V_\ell$ (the unbiased case can be considered analogously), where V_ℓ is the estimator variance on the ℓ^{th} level, of the summary statistic of interest (here the first moment $\mathbb{E}(X_i)$) to within a user prescribed tolerance, ϵ .

If a pair of paths on the ℓ^{th} level takes c_ℓ units of CPU time to be generated, we wish to minimise

$$\sum_{\ell=0}^L n_\ell c_\ell, \quad (19)$$

subject to

$$\sum_{\ell=0}^L V_\ell < \epsilon. \quad (20)$$

This optimisation problem can be solved using the method of Lagrange multipliers¹³ and suggests that we should take

$$n_\ell = \left\lceil \frac{V_\ell/c_\ell}{\epsilon \sum_{\ell=0}^L \sqrt{c_\ell V_\ell}} \right\rceil, \quad (21)$$

where $\lceil \cdot \rceil$ represents the ceiling function. One means to estimate c_ℓ and V_ℓ involves generating an initial number of sample paths on each level (usually 100 or 1000)¹.

III. RECONSTRUCTING CUMULATIVE DISTRIBUTION FUNCTIONS

In this section we introduce two methods to approximate CDFs of discrete state systems and then how to implement these methods efficiently alongside the multi-level method. We demonstrate the utility of our approach using two explanatory examples.

A. Using a maximum entropy approach

We first consider approximating the CDF of species X_i at terminal time T using a finite set of M moments, $\{\mu_1, \dots, \mu_M\}$, and the method of maximising entropy¹¹ to first estimate the corresponding probability distribution function (PDF). For a distribution on the discrete

state space $\Omega = \{x_1, \dots, x_m\}$ with PDF $\mathbf{p} = (p_1, \dots, p_m)$, the entropy function, $\mathbb{H}(\mathbf{p})$, is defined as

$$\mathbb{H}(\mathbf{p}) := \sum_{i=1}^m p_i \log \left(\frac{1}{p_i} \right). \quad (22)$$

If $m = \infty$ we select an integer N at which to truncate the state space, so that we work with $\Omega_t = \{x_1, \dots, x_N\}$. The optimisation problem is now to find the PDF, \mathbf{p} , which maximises the entropy function (22) subject to the moment constraints

$$\mu_j = \sum_{i=1}^N (x_i)^j p_i, \quad (23)$$

for $j \in \{0, \dots, M\}$. Note we include μ_0 , and set $\mu_0 = 1$ as we require $\sum_{i=1}^N p_i = 1$. The solution is obtained by introducing $M + 1$ Lagrange multipliers¹⁸, $\boldsymbol{\lambda} = (\lambda_0, \dots, \lambda_M)$, and the Lagrange functional

$$\Gamma = \sum_{i=1}^N p_i \log \left(\frac{1}{p_i} \right) - \sum_{j=1}^M \lambda_j \left(\sum_{i=1}^N (x_i)^j p_i - \mu_j \right) - c \left(\sum_{i=1}^N p_i - \mu_0 \right), \quad (24)$$

where $c = \lambda_0 - 1$. By solving $\partial\Gamma/\partial p_i = 0$, we obtain the solution

$$p_i = \exp \left(- \sum_{\ell=0}^M \lambda_\ell (x_i)^\ell \right), \quad (25)$$

for $i \in \{1, \dots, N\}$. Now that we know the form of the solution, which has $M + 1$ unknowns, $\boldsymbol{\lambda} = (\lambda_0, \dots, \lambda_M)$, we have a system of $M + 1$ nonlinear equations,


$$\sum_{i=1}^N \exp \left(- \sum_{\ell=0}^M \lambda_\ell (x_i)^\ell \right) (x_i)^j = \mu_j, \quad (26)$$

for $j \in \{0, \dots, M\}$. We need to employ a numerical method to calculate $\boldsymbol{\lambda} = (\lambda_0, \dots, \lambda_M)$, the one used in this paper is adapted from work by Mohammad-Djafari¹⁵. The CDF can then be trivially calculated.

We now present possible ways of efficiently implementing this method in combination with the multi-level method in order to estimate CDFs of given species numbers. To aid us we use two example models, the Schlögl model and a dimerisation model.

1. The Schlögl model

The Schlögl model^{16,21} consists of a single species, A , and four possible reactions:



We fix the rates to be $[k_1, k_2, k_3, k_4] = [2200, 37.5, 0.18, 0.00025]$ for the remainder of this paper. We chose the initial condition $A(0) = 0$ and terminal time $T = 20$. The first question we need to answer is how many moments are necessary to approximate the CDF of A to within a desired degree of accuracy. To gain insight into this problem, we generated 10^6 sample paths using Gillespie's DM and calculated the first seven moments. With these moments we calculated approximate PDFs for $A(20)$ using up to and including seven moments in the method of maximum entropy. The results are presented on the left of Figure 1.




FIG. 1. Left: PDFs from the maximum entropy method using up to the first seven moments on the truncated state space $\Omega_t = \{0, 1, 2, \dots, 600\}$. Right: empirical PDF and the PDF generated using the maximum entropy method and the first seven moments.

Figure 1 indicates that the peaks of the PDFs change significantly with the addition of the fourth, fifth and sixth moments into the entropy calculation. Thus it appears that for the Schlögl system we need at least six moments to accurately approximate the CDF of species A . The diagram on the right of Figure 1 compares the empirical PDF and the approximate PDF using seven moments. Qualitatively, the approximate PDF captures the behaviour of the empirical PDF very well. However, in order to have a quantitative measure of the accuracy of approximate PDFs generated with different numbers of moments, we use the Kolmogorov-Smirnov distance¹⁷.

Given two cumulative distribution functions (CDFs), $F(x)$ and $G(x)$, defined on support Ω , the Kolmogorov-Smirnov distance is given by

$$\mathbb{D}_{F,G} := \sup_{x \in \Omega} |F(x) - G(x)|. \quad (28)$$

The Kolmogorov-Smirnov distance can be interpreted as the largest vertical distance between the two CDFs. The Kolmogorov-Smirnov distance between the approximate distribution calculated using the first seven moments and the empirical distribution is 0.0037. As a comparison we define $c_i = \mathbb{E}[\mathbf{1}\{A(20) \leq i\}] = \mathbb{P}(A(20) \leq i)$ for $i \in \Omega$, and calculate the 95% confidence intervals of each estimate c_i , the largest of which was 0.0001. This means that the source of the error is in the maximum entropy method and is not a consequence of noise in the empirical CDF.

The multi-level method is used to generate the moments required by the maximum entropy method. In practice we do not want to rely on results from an eSSA to calculate the Kolmogorov-Smirnov distance so instead we assume that the PDFs on the left of Figure 1 converge to the empirical distribution. We can therefore choose the number of moments, i , to be the minimal i such that $\mathbb{D}_{P_i, P_{i-1}} < \delta$, where δ is chosen to provide the desired level of accuracy and P_i represents the CDF calculated using the first i moments in the method of maximum entropy.

The question we now need to answer is which moment should have its estimator variance controlled when choosing the number of paths on each level in the multi-level method. Recall in Section II we controlled the estimator variance of the first moment as this was the only system statistic of interest. However, now we have multiple system statistics. We choose to control the accuracy of every moment, $\mathbb{E}(X^k)$, by controlling its variance, $\text{Var}(X^k)$, to within $\epsilon \cdot \mathbb{E}(X^k)$ for some $\epsilon > 0$. We now present an adaptive algorithm to generate an approximate CDF using the multi-level method together with the method of maximum entropy:

1. Set the number of levels L and the number of initial moments $m \geq 2$. Set the variance control parameter, ϵ , and the tolerance parameter for the Kolmogorov-Smirnov distance test, δ .
2. Generate 1000 initial sample paths on each level using the multi-level algorithm.
3. Calculate the number of additional paths, n_ℓ^j , necessary on each level, ℓ , to ensure $\text{Var}(X^j) < \epsilon \cdot \mathbb{E}(X^j)$ using equation (21) for each $j \in \{1, \dots, m\}$. Find $n_\ell^{k_\ell} = \max_{i \in \{1, \dots, m\}} \{n_\ell^i\}$
4. Generate $n_\ell^{k_\ell}$ sample paths on each level, ℓ , using the multi-level algorithm.

5. Let P_1 be the CDF generated using the maximum entropy method using $m-1$ moments, and P_2 the CDF using m moments. If $\mathbb{D}_{P_1, P_2} < \delta$ terminate the algorithm, otherwise let $m = m + 1$ and return to step 3.

This adaptive algorithm systematically increases the number of moments used in the entropy calculation until the addition of a new moment has negligible effect on the resulting CDF, which is determined by the tolerance parameter, δ . The algorithm relies on a convergence assumption, i.e. increasing the number of moments indefinitely will yield CDFs that converge to the true CDF. We note that, however, in practice, the algorithm is limited by the robustness of the numerical scheme chosen to calculate the CDFs.

In Figure 2 we present for the Schlögl model: the CDF, $P(x)$, generated using the adaptive algorithm with $\epsilon = 0.03$ and $\delta = 0.02$; and the empirical CDF, $F(x)$, generated using Gillespie's DM which took 4568 seconds and 38853 seconds of CPU time to compute, respectively. The adaptive algorithm selected six as the optimal number of moments.





FIG. 2. Left: The blue curve is the CDF $F(x)$ generated using Gillespie's DM with 10^6 sample paths. The red curve is the CDF $P(x)$ generated using the adaptive multi-level algorithm and the method of maximum entropy, with control parameters $\epsilon = 0.03$ and $\delta = 0.02$. The truncated state space is $\Omega_t = \{0, 1, 2, \dots, 600\}$. The Kolmogorov-Smirnov distance is 0.0125. Right: A plot of the error: $|F(x) - P(x)|$.

2. A dimerisation model

As our next example we consider gene expression for a single gene¹ (G). It transcribes mRNA, known as messengers (M), which can, in turn, translate proteins (P). Pairs of these proteins may combine to form a dimer (D). The mRNA and proteins may also degrade. This biochemical network is described by the following set of reactions¹:



We fix the rates to be $[k_1, k_2, k_3, k_4, k_5] = [25, 1000, 0.001, 0.1, 1]$ for the remainder of this paper. We wish to approximate the distribution of the dimer population at a fixed terminal time $T = 1$ and we take the initial conditions to be $[G(0), M(0), P(0), D(0)]^T = [1, 0, 0, 0]^T$.




FIG. 3. Left: The blue curve is the CDF $F(x)$ generated using Gillespie's DM with 10^6 sample paths. The red curve is the CDF $P(x)$ generated using the adaptive multi-level algorithm and the method of maximum entropy, with control parameters $\epsilon = 0.0025$ and $\delta = 0.01$. The truncated state space is $\Omega_t = \{0, 1, 2, \dots, 10000\}$. The Kolmogorov-Smirnov distance is 0.0029. Right: A plot of the error: $|F(x) - P(x)|$.

Figure 3 shows the CDF, $P(x)$, produced using the adaptive multi-level algorithm, it took 321 seconds of CPU time to compute. The figure also contains the empirical CDF, $F(x)$, generated using Gillespie's DM, which took 2598 seconds of CPU time to compute. The maximum 95% confidence interval for the estimates c_i , where $c_i = \mathbb{E}[\mathbf{1}\{D(1) \leq i\}]$ for $i \in \Omega$, calculated using 10^6 realisations of Gillespie's DM was 0.0001. This confirms that the error

lies in the maximum entropy method and is not a consequence of noise in the empirical distribution. The adaptive algorithm selected four as the optimal number of moments.

Figure 4 shows how the CPU times increase for both the Schlögl model and the dimerisation model as the variance control parameter, ϵ , is decreased. The CPU times for the adaptive algorithm using the method of maximum entropy for the Schlögl model are seen to jump at around $\epsilon = 0.02$. This is because as ϵ is decreased the number of moments necessary to pass the Kolmogorov-Smirnov distance test increases, which results in a significant increase in the CPU time.




FIG. 4. Left: CPU times for the adaptive algorithm using the method of maximum entropy for the Schlögl model where $\delta = 0.02$. Right: CPU times for the adaptive algorithm using the method of maximum entropy for the dimerisation model where $\delta = 0.01$. In both plots the red line shows the CPU for 10^6 sample paths from Gillespie’s DM.

B. Using an indicator function approach

We now consider another method to approximate CDFs. Let X be a distribution on a discrete state space Ω . We can approximate the CDF of X directly by estimating the following statistics,

$$c_i = \mathbb{P}(X \leq i) = \mathbb{E}[\mathbf{1}\{X \leq i\}], \quad (30)$$

for $i \in \Omega$. Although we refer to the collection of estimates c_i , for $i \in \Omega$, as an approximate CDF it is important to note that this method does not always generate a true CDF because estimates are not necessarily monotonic increasing i.e. $c_i \leq c_{i+1}$ for $\{i, i+1\} \in \Omega$. For two

solutions to this problem see Section IV.

These statistics, c_i , can be estimated using the multi-level method for which we control the variance of each of the c_i , requiring, as before, $\text{Var}(\mathbf{1}\{X \leq i\}) < \epsilon$ for some $\epsilon > 0$. One option is then to generate $N_\ell = \max_{i \in \Omega} \{n_\ell^i\}$ paths on each level, where n_ℓ^i is the number of paths required on the ℓ^{th} level to control the variance of c_i . However, this requires large CPU times. Instead we introduce I where $I \subset \Omega$. We generate $\hat{N}_\ell = \max_{i \in I} \{n_\ell^i\}$ on each level and then calculate the estimates, c_i , for $i \in \Omega$. Then we iteratively increase the size of the subset I until further increases have only negligible effects on the approximate CDFs produced. We update I by finding the maximal set $J \subset I$ where each element $j \in J$ corresponds to the estimate c_j that required the largest number of paths on some level, ℓ , of the multi-level method. Then, for each $j \in J$, we add new entries to I as the midpoints of j and the nearest neighbours of j that are already in I . We now present an algorithm for the adaptive method and the indicator function approach:

1. Set the number of levels L and the variance tolerance parameter, ϵ , for all the estimates c_i . Set the tolerance parameter, δ , for the Kolmogorov-Smirnov distance test. Set the initial subset $I \subset \Omega$ from which the number of paths to be generated will be decided.
2. Generate 1000 initial sample paths on each level using the multi-level method.
3. Calculate the number of additional paths, n_ℓ^i , necessary on each level, ℓ , to ensure $\text{Var}(\mathbf{1}\{X \leq i\}) < \epsilon$ for each $i \in I$ using equation (21).
4. Generate \hat{N}_ℓ paths on each level, ℓ , where \hat{N}_ℓ is given by,

$$\hat{N}_\ell = \max_{i \in I} \{n_\ell^i\}, \quad (31)$$

using the multi-level method.

5. Let P_1 be the approximate CDF generated by estimating c_i for each $i \in \Omega$.
6. Construct the maximal set J such that for every $j \in J$, $n_\ell^j = \hat{N}_\ell$ for some level ℓ . Then for every $j \in J$ add additional entries to I as the midpoints between j and its direct neighbours. If I cannot be updated further, i.e. there is no midpoint not already in I , then terminate the algorithm.

7. Calculate the number of additional paths, n_ℓ^i , necessary on each level ℓ to ensure $\text{Var}(\mathbf{1}\{X \leq i\}) < \epsilon$, for each $i \in I$ using equation (21).
8. Generate \hat{N}_ℓ paths on each level, ℓ , where \hat{N}_ℓ is given by,

$$\hat{N}_\ell = \max_{i \in I} \{n_\ell^i\}, \quad (32)$$

using the multi-level method.

9. Let P_2 be the approximate CDF created by calculating c_i for each $i \in \Omega$.
10. If $\mathbb{D}_{P_1, P_2} < \delta$, terminate the algorithm, otherwise let $P_1 = P_2$ and return to step 6.

We note that $c_i \in [0, 1]$ for every $i \in \Omega$ and so we do not scale the variance tolerance of each estimate by c_i , instead we select a single variance tolerance, ϵ , for all estimates. We now use this adaptive algorithm on the two example systems we have seen previously.

1. The Schlögl model

We consider the molecular species A and the reactions (27). We truncate the state space to $\Omega_t = \{0, \dots, 600\}$ and consider the initial condition $A(0) = 0$ and terminal time $T = 20$. The initial subspace used is $I = \{100, 200, 300, 400, 500\}$. In Figure 5 we present the approximate CDF $P(x)$ generated using the adaptive multi-level algorithm with $\epsilon = 0.0025$ and $\delta = 0.001$ and the CDF $F(x)$ generated using Gillespie's DM which took 6762 seconds and 38853 seconds of CPU time, respectively. The CDFs are very similar and the exact error is given on the right of Figure 5.

2. A dimerisation model

Consider the four species G , M , P and D and the reactions (29). We wish to approximate the distribution of the dimer population at a fixed time $T = 1$ where we take the initial conditions to be $[G(0), M(0), P(0), D(0)]^T = [1, 0, 0, 0]^T$. We truncate the state space to $\Omega_t = \{0, \dots, 10000\}$ and the initial subspace is $I = \{1000, 2000, \dots, 8000, 9000\}$. In Figure 6 we present the approximate CDF $P(x)$ generated using the adaptive multi-level algorithm with $\epsilon = 0.002$ and $\delta = 0.001$ and the CDF $F(x)$ generated using Gillespie's DM which took




FIG. 5. Left: The blue curve is the CDF $F(x)$ generated using Gillespie’s DM with 10^6 sample paths. The red curve is the approximate CDF $P(x)$ generated using the adaptive multi-level algorithm for the indicator function approach, with control parameters $\epsilon = 0.0025$ and $\delta = 0.001$. The truncated state space is $\Omega_t = \{0, 1, 2, \dots, 600\}$. The Kolmogorov-Smirnov distance is 0.0042. Right: A plot of the error: $|F(x) - P(x)|$.

249 seconds and 2598 seconds of CPU time, respectively. Again the CDFs are very similar and the exact error is given on the right of Figure 6.

Figure 7 shows how the CPU times increase for both the Schlögl model and the dimerisation model as the variance control parameter, ϵ , is decreased. As ϵ is decreased we see that the CPU time of the adaptive method and the indicator function approach becomes larger than the CPU time for Gillespie’s DM.

IV. DISCUSSION

We have presented two novel algorithms designed to employ the multi-level Monte Carlo method to generate approximate CDFs for molecular populations at a terminal time T . The first was the method of maximum entropy which used the multi-level method to generate a finite set of moments and from these moments calculated an approximate distribution. We note that this method is favourable if an analytical approximation of the PDF is required. The second method was an indicator function approach. This adaptive method focussed on generating more paths to control the estimates of c_i in regions of Ω where the random variables $\mathbf{1}\{X \leq i\}$ have a higher variance. These regions are often found to be near the




FIG. 6. Left: The blue curve is the CDF $F(x)$ generated using Gillespie’s DM with 10^6 sample paths. The red curve is the approximate CDF $P(x)$ generated using the adaptive multi-level algorithm for the indicator function approach, with control parameters $\epsilon = 0.002$ and $\delta = 0.001$. The truncated state space is $\Omega_t = \{0, 1, 2, \dots, 10000\}$. The Kolmogorov-Smirnov distance is 0.0028. Right: A plot of the error: $|F(x) - P(x)|$.

peaks of the corresponding PDFs which is where the method of maximum entropy returns the largest errors. Thus we see that the indicator function approach achieves smaller Kolmogorov-Smirnov distances than the method of maximum entropy when using the Schlögl model.

However, we note that the indicator function approach does not necessarily provide a CDF as the estimates c_i are not always monotonic increasing in i . Although the estimates can still be used to give an accurate and meaningful description of the distribution, if a physical CDF is required there are at least two solutions. Firstly one can vastly decrease the value of the variance control parameter ϵ to reduce the probability of non-monotonicity, this however will significantly increase the CPU time of the method such that it is comparable with the tau-leap method and Gillespie’s DM. Secondly one can take estimates for c_i from a coarse subset $S \subset \Omega$, and generate the CDF using polynomial interpolation whilst enforcing a positive gradient over the range of Ω . This second approach uses a coarser subset S to reduce the probability of non-monotonicity without the further increase in the CPU time.

Future work will involve investigating the potential for improvements on the efficiency of the maximum entropy method, as well as improvements in the efficiency of using multi-level methods to produce approximate CDFs in general.




FIG. 7. Left: CPU times for the adaptive algorithm using the indicator function approach for the Schlögl model where $\delta = 0.001$. Right: CPU times for the adaptive algorithm using the indicator function approach for the dimerisation model where $\delta = 0.001$. In both plots the red line shows the CPU for 10^6 sample paths from Gillespie’s DM.

ACKNOWLEDGEMENTS

The authors would like to thank Chris Lester for many insightful discussions.

REFERENCES

- ¹D. F. Anderson and D. J. Higham. Multi-level Monte Carlo for continuous time Markov chains with applications in biochemical kinetics. *SIAM Multiscale Modeling and Simulation*, 10(1):146–179, 2012.
- ²M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- ³S. Engblom. Spectral approximation of solutions to the chemical master equation. *Journal of Computational and Applied Mathematics*, 229(1):208:221, 2009.
- ⁴R. Erban, J. S. Chapman, I. G. Kevrekidis, and T. Vejchodský. Analysis of a stochastic chemical system close to a sniper bifurcation of its mean-field model. *SIAM Journal of Applied Mathematics*, 70(3):984:1016, 2009.
- ⁵R. Erban, J. S. Chapman, and P. K. Maini. A practical guide to stochastic simulations of reaction-diffusion processes. *arXiv*, 0704(1908), 2007.

- ⁶M. B. Giles. Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617, 2008.
- ⁷D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- ⁸D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- ⁹T. Jahnke. On reduced models for the chemical master equation. *Multiscale Modelling and Simulation*, 9(4):1646–1676, 2011.
- ¹⁰T. Jahnke and W. Huisinga. Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology*, 54(1):1–26, 2007.
- ¹¹E. T. Jaynes. *Probability Theory: The Logic of Science*. Cambridge University Press, 2003.
- ¹²T. G. Kurtz and D. F. Anderson. Continuous time Markov chain models for chemical reaction networks. In *Design and Analysis of Biomolecular Circuits*, pages 3–42. Springer, 2011.
- ¹³C. Lester, R. E. Baker, M. B. Giles, and C. A. Yates. A guide to efficient discrete-state multi-level simulation of stochastic biological systems. *arXiv*, 1412(4069), 2014.
- ¹⁴C. Lester, C. A. Yates, M. B. Giles, and R. E. Baker. An adaptive multi-level simulation algorithm for stochastic biological systems. *The Journal of Chemical Physics*, 142(2):024113, 2015.
- ¹⁵A. Mohammad-Djafari. A Matlab program to calculate maximum entropy distributions. *Fundamental Theories of Physics*, 50:221–233, 1992.
- ¹⁶F. Schlögl. Chemical reaction models for non-equilibrium phase transitions. *Zeitschrift für Physik*, 253(2):147:161, 1972.
- ¹⁷D. J. Sheskin. *Handbook of Parametric and Non-parametric Statistical Procedures*. Chapman and Hall, third edition, 2003.
- ¹⁸J. Stewart. *Calculus*. Brooks/Cole, sixth edition, 2009.
- ¹⁹E. Süli and D. F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003.
- ²⁰T. Székely, K. Burrage, and R. Erban. A higher order numerical framework for stochastic simulation of chemical reaction systems. *BMC Systems Biology*, 6(1):85, 2012.

²¹T. Wilhelm. The smallest chemical reaction system with bistability. *BMC Systems Biology*, 3(1):90, 2009.